

Education Working Paper Archive

Collective Bargaining and the Performance of the Public Schools

May 25, 2007

Terry M. Moe

Professor, Department of Political Science

Stanford University

and

Senior Fellow, Hoover Institution

Abstract

Students of American politics have had little to say about public sector unions and their impacts on government. There is, of course, a vast literature on public bureaucracy. But that literature has always emphasized that the power of bureaucrats is rooted in their expertise, or in the entrepreneurial activities of agency leaders. It has largely ignored the fact that bureaucrats can and often do join unions to promote their own interests, and that the power of these unions may have important consequences for—and help explain—the policies, organization, and performance of government. In this paper, I study the effects of unionization on public education. Specifically, my focus here is on the public schools, which are among the most common form of government agency in the United States, and I investigate whether collective bargaining by teachers—the key bureaucrats in this case—affects the capacity of the schools to educate children. Using data from the state of California, the analysis shows that, in large school districts, the restrictiveness of the teacher contract has a very negative impact on academic achievement. It also shows that, in these large districts, restrictive contract rules are especially negative in their effects on the academic achievement of minorities. The evidence presented here, then, suggests that public sector unions do indeed have important consequences for American public education. Whether they are consequential in other areas of government remains to be seen, but it is an avenue of research well worth pursuing.

Collective Bargaining and the Performance of the Public Schools

Prior to 1960, few public workers in the United States belonged to unions and almost none were covered by collective bargaining. All this changed during the next two decades, as most states passed laws making it easier for public sector unions to organize and gain bargaining rights. The result was a period of explosive growth in which the portion of public workers covered by collective bargaining grew to more than 40% by the early 1980s. In the aggregate, it has remained at roughly that level ever since, and is actually much higher in many states, localities, and occupations, particularly outside the right-to-work South. Meantime, unions in the private sector have been caught in a disastrous free-fall, with coverage at just 8.5% as of 2005 (U.S. Census Bureau, 2007).

Public sector unions have changed the dynamics of American politics. They have compelling incentives to be politically active, because their members are directly dependent on government for their livelihoods. And they have acted on these incentives with seriousness and dedication, using impressive reservoirs of money and manpower to make themselves formidable players in electoral campaigns, as well as in every major sphere of governmental decision making: legislative, administrative, and judicial (Moe, 2006a; Johnson and Libecap, 1994; Blasé, Blake, and Dion, 1997; Troy, 1994).

They have done more, however, than change the nation's politics. For when they engage in collective bargaining, the contracts that result—each of which may impose hundreds of formal rules—become integral components of the structure of government. They are often major players, then, in determining how government is organized. State and local governments, where most of the nation's policies are carried out and most of its money spent, are more affected by collective bargaining than the federal government. And some state and local governments are more affected than others. But such variation simply makes the phenomenon more interesting and worth studying. On the whole, collective bargaining is quite common within the public sector, it is often a source of governmental structure,

and—if institutions really do matter—it is likely to be an important (if varying) influence on governmental performance (Freeman, 1986; Freeman and Ichniowski, 1988; Lewin et al., 1988).

Students of American politics have had little to say about the rise of public sector unions, and little to say, in particular, about the impact of collective bargaining on the structure and performance of government.¹ There is, of course, a vast literature on public bureaucracy. But that literature—from the early writings of Weber (1947) to the most recent developments in rational choice modeling (e.g., Huber and Shipan, 2002; Epstein and O'Halloran, 1999)—has always emphasized that the power of bureaucrats is rooted in their expertise: in the knowledge and experience that give them leverage in dealing with their political superiors (and with clients). It has also recognized the various ways in which bureaucratic leaders can establish reputations, build constituencies and coalitions, and otherwise take entrepreneurial action to gain power and autonomy in politics (Carpenter, 2001; Rourke, 1984).

While these aspects of bureaucratic power are surely important, the literature largely ignores the fact that even the most ordinary bureaucrats can get organized into unions. When they do, their numbers and resources may well translate into considerable power over the policies, structures, and performance of government, power that is used to promote their own occupational interests—in material benefits, in job security, in restrictive work rules, in public policy—and that may lead to a host of impacts on government. Many of these impacts are likely to be intentional, as when unions succeed in blocking unwanted policy reforms or in securing new work rights. Some impacts, on the other hand, may be quite unintentional—as could occur, for instance, if restrictive work rules have the effect of making public agencies less productive. But either way, they are products of union power, and they stand to have important consequences for what government is and does.

¹ The major exception is Johnson and Libecap (1994). This is a book about federal employee unions, however, and most of the membership and power of U.S. public sector unions are at the state and local levels. See also Carpenter's (2001) analysis of the activities of the early postal unions.

Students of American politics, including students of bureaucracy, have long recognized the power of special interest groups, including unions. But the groups they study are outside of government—and public sector unions arise from the inside, representing the special interests of the government's own employees.² It is time for political scientists to make public sector unions a serious part of their agenda for theory and research.

This paper is a step in that direction. My focus here is on the public schools, which are among the most common form of government agency in the United States and surely among the most important. The key bureaucrats in this case are the schools' teachers, who, outside the South, tend to be heavily unionized and covered by collective bargaining. The basic question here is whether, in using their power to secure contract rules that advance the occupational interests of their members, the teachers unions are (unintentionally) limiting the capacity of the public schools to educate children.

The data are from a sample of schools and districts from the state of California. Collective bargaining contracts are coded in terms of how restrictive they are in limiting control by superiors and imposing a structure favorable to the occupational interests of teachers. The analysis then shows that, in large school districts, the restrictiveness of the teacher contract has a very negative impact on school performance: the more restrictive the contract, the less successful are the public schools at promoting student achievement. It also reveals that, in these large districts, restrictive contract rules are especially negative in their effects on the academic achievement of minorities.

The evidence presented here, then, suggests that public sector unions do indeed have important consequences for the performance of the public schools. Whether they are consequential in other areas of government remains to be seen, but it is an avenue of research well worth pursuing.

² The interest group literature, like the bureaucracy literature, has very little to say about public sector unions. See, e.g., Baumgartner and Leech, 1998; Ciglar and Loomis, 2006.

Expectations

Since the publication of A Nation at Risk (National Commission on Excellence in Education, 1983), the watershed report that warned of a “rising tide of mediocrity” in America’s schools, public officials have been under intense pressure to improve academic performance. They have responded with billions of additional dollars and countless reforms, and with a commitment to school improvement that, remarkably, has maintained its salience and drive throughout the last quarter century. For the most part, the focus has been on promoting academic achievement. But there has been considerable concern as well for closing the “achievement gap” between white and minority children, and for improving the large urban districts that many minority students attend (Peterson, 2003).

On both counts, the most fundamental changes have come about through new accountability systems, which, through the imposition of more rigorous standards, testing regimes, and consequences attached to performance, represent the most aggressive effort yet by public authorities to improve academic achievement. Accountability reforms spread rapidly across the states during the 1990s, and in 2001 the federal government enacted No Child Left Behind, landmark legislation that imposed accountability rules for the country as a whole: rules designed to promote academic improvement, but also to spotlight the achievement of minority students, and to insist on efforts to close the achievement gap (Peterson and West, 2003).

Accountability reforms are clear indications—obvious to administrators, teachers, and students at the school level every day—that public officials are serious and in broad agreement about increasing student learning. Behavior in the schools is shaped by other structures as well, however. And notable among them are the rules imposed by collective bargaining. What should we expect of these contract rules? Do they contribute to the authorities’ efforts to improve academic achievement? Or do they tend to get in the way?

There is good reason to think that the latter is much more likely, at least on balance. Collective bargaining would not exist except for the power of the teachers unions, and they are not in the business of promoting academic achievement. They are in the business of pursuing their own interests, which arise from their survival-based concerns for maintaining membership and financial resources, as well as their induced concern (because it contributes to these more basic objectives) for representing the occupational wants and needs of their members. They have an interest, for example, in protecting member jobs and fighting for higher wages and fringe benefits. They also have an interest in fighting for greater teacher autonomy, less threatening methods of evaluation, prohibitions on non-classroom duties, smaller classes, fewer course preparations, and similar protections and benefits that teachers value. The unions secure these objectives through formal contract rules that require or prohibit certain behaviors on the part of management—and, most generally, place restrictions on top-down control (e.g., McDonnell and Pascall, 1979; Grimshaw, 1979; Hoxby, 1996; Moe, 2006b).

These restrictions ensure that the public schools are literally not organized to promote academic achievement. When contract rules make it difficult or impossible to weed out mediocre teachers, for example, they directly undermine the single most important determinant of student learning: teacher quality (Sanders and Rivers, 1996). And when contract rules guarantee teachers seniority-based transfer rights, they ensure that teachers cannot be allocated to their most productive uses (Levin, Mulhern, and Schunck, 2005). Much the same can be said about a long list of standard contract provisions: they often come into conflict with basic requisites of effective organization. This is precisely what we should expect. The contract rules were never intended to make the schools effective.

Still, there are shades of gray here. For when teachers unions pursue their own interests, they will sometimes take actions that happen to be good for student achievement too. Their pursuit of smaller classes is an obvious example. More generally, because unions tend to secure better wages,

benefits, and working conditions, and because they expand worker autonomy in the workplace, they may help to attract higher quality workers, promote professionalism, and lessen worker absenteeism and turnover, and thus to enhance productivity (Freeman and Medoff, 1984).

These are plausible arguments. Yet in the case of the teachers unions there are persuasive reasons for thinking that the positives are likely to be small relative to the negatives. One reason is that any positive effects are accidental by-products of what the teachers unions do in their own interests. They pressure for smaller classes, for instance, because teachers like them and because they call for more hiring (and more union members), not because they are good for student achievement. If there were an optimal class size for student achievement (given district budget constraints and opportunity costs), this would not stop the unions from pressuring for class-sizes that are still smaller. The same is true for seniority rights. Unions pressure for them because they take job assignments out of the hands of administrators, not because they increase experience or reduce turnover; and if the lack of managerial control over the allocation of teachers has negative consequences for student achievement that outweigh their impacts on experience and turnover, the unions would continue to support the seniority system anyway. The same logic applies to issues of professionalism and other aspects of union-inspired organization. The bottom line is that the interests of teachers (and unions) are simply not aligned with the interests of children, and the organizational arrangements pursued by unions will ultimately diverge from those that are best for students.

We also need to recognize that any contract provisions that happen to be good for student achievement could be adopted (as policies) by school districts on their own. And especially in this era of state and federal accountability pressures, they would have incentives to do just that. They would not, on the other hand, have incentives to adopt contract provisions that have negative effects on student achievement—but would adopt them only because the unions are powerful and demand them.

The central hypothesis I will be testing in this paper, then, can be stated as follows: the restrictiveness of the collective bargaining contract—its overall limitation on managerial control in the interests of teachers (and unions)—has negative consequences, on balance, for the performance of the public schools. The more restrictive the contract, the more difficult it will be for schools to do their jobs well, and the less students will learn.

This hypothesis—and this paper—are just starting points. A more detailed investigation of collective bargaining would suggest that, while there are persuasive reasons for expecting these labor contracts to have negative effects, there are also reasons for thinking that the magnitudes of the effects are not everywhere the same, but are likely to vary with a range of conditions. Taking on this larger task would require a far more extensive analysis and is beyond the scope of this one paper. But once I have presented the basic empirical findings on restrictiveness, I will go on to discuss two conditions of special concern in this era of education reform—the size of the district and the concentration of minority students—and show that they do indeed appear relevant to how collective bargaining affects the schools.

Research on the Impact of Teachers Unions

There is a rather large research literature on the impact of unions. These studies generally agree that unionization leads to higher costs in both the private and public sectors. But the findings are mixed on the question of how unions affect overall productivity, and thus whether the higher costs are counterbalanced by increases in output. One reason for the mixed findings is that these studies often use different methods and measures, are carried out on widely varying industries and economic contexts, and are not of equal quality, all of which make summary conclusions difficult. It is perhaps not surprising, then, that even after decades of research the dispute about the overall impact of unions lives on without resolution (e.g., Doucouliagos and Laroche, 2003; Hirsch, 2004).

A small subset of this research has dealt specifically with teachers unions, attempting to determine whether collective bargaining within the public school system has an impact on student learning. Here too the findings have been mixed. I will not review these studies individually, but a brief look at some of their characteristics helps to suggest how extensive the heterogeneity is, and why it can easily lead to disparate and unreliable findings (Eberts and Stone, 1984; 1986; Milkman, 1997; Kurth, 1987; Grimes and Register, 1990; Argyris and Rees, 1995; Peltzman, 1993; Hoxby, 1996; Nelson and Rosen, 1996; Steelman, Powell, and Carini, 2000).

(1) Many studies are carried out at the state level, and are thus based on data so heavily aggregated—average achievement scores for entire states, for example—that there is little hope of discovering causes and effects at the district level, where collective bargaining actually takes place.

(2) Collective bargaining is measured differently in different studies. Some use dummy variables to indicate whether a state or district has collective bargaining or not. Others focus instead on union strength, using the percentage of teachers belonging to unions as a proxy.

(3) Achievement is typically measured in one of two ways: mathematics test scores or scores on SAT / ACT exams. Yet SAT and ACT tests are only taken by college-bound seniors, who are unrepresentative of students generally. And math scores capture just a small part of the academic curriculum and are narrow measures of achievement. The populations whose math scores are featured, moreover, vary dramatically from study to study: Eberts and Stone (1984, 1986), for example, focus on 4th graders, Argyris and Rees (1995) target 10th graders, and Milkman (1997) studies 12th graders.

(4) Some studies are based on student-level data, and thus on samples that are quite large. Eberts and Stone (1984, 1986), for instance, analyze a sample of approximately 14,000 students in 328 elementary schools, and other studies have samples in the thousands as well (Argyris and Rees, 1995; Milkman, 1997; Grimes and Register, 1990). Ordinarily, large samples would be a good thing. But in

these studies, many students are drawn from the same schools and districts; and because of this clustering, they should not be treated as independent observations. Were clustering properly taken into account, standard errors would tend to be higher and statistical significance more difficult to achieve. Yet none of these studies adjusts for clustering. If they did, their findings about the impact of collective bargaining would likely be weakened.

(5) States with low union density or little or no collective bargaining are almost all southern or border states, and the school systems in these states have historically been low performers. This simple fact makes it difficult to disentangle regional influences from union influences on student achievement. It also raises endogeneity issues, because some of the factors that explain why unions are weak in these states may also explain why the school systems do not perform well; their political cultures, for instance, generate attitudes and laws hostile to unions, but also low spending, low taxing, and low public pressure for good schools. Many of the studies (Hoxby, 1996, is the exception) ignore these problems entirely, while others address it by merely introducing a dummy for the South.³

In view of how different these studies are and how many questions are raised by their methods, there can be little surprise that this literature has not led to a coherent set of findings. One study, however, stands apart from the others and is due special weight. This is the analysis by Hoxby (1996), which assesses the impact of collective bargaining by looking at districts before and after the unions gain bargaining rights, a unique design that only she has employed. What she shows, in the literature's most sophisticated analysis, is that collective bargaining increases school inputs—total spending, teacher salaries, teacher-student ratios—but also decreases their productivity, so that the unions' overall impact on school performance is actually negative (as measured by the drop-out rate).

³ Endogeneity issues may also arise for other reasons, which will be discussed at a later point in the text.

A New Study

My own approach departs from the others in this literature in two basic respects. First, I am concerned more narrowly with the restrictiveness of the contract rather than with collective bargaining or union strength per se. And second, I am comparing its effects across jurisdictions that all engage in collective bargaining and whose teachers are virtually all unionized.

The focus on restrictiveness means that, rather than reducing collective bargaining to gross dichotomies (have it / don't have it) or gross proxies for union strength (percentage unionized), we are led to explore the actual contents of the labor contracts themselves, and thus to determine what rules they contain and what implications they have for student achievement. In this way, we can gain a new and more informative angle on the academic impact of collective bargaining.

By looking only at districts that have collective bargaining, we gain additional advantages. The vast majority of school districts (of any size) in this country do engage in collective bargaining. This is business as usual in public education. An approach that compares these “normal” districts to those that do not have collective bargaining is a risky proposition; for as I suggested above, there are reasons the latter districts have never gotten unionized, and these factors (if not measured or otherwise adjusted for) may lead to biases. By looking just at the “normal,” unionized districts, we are looking at districts that have a great deal in common, and the risk of bias should be reduced considerably.

In addition to these basic differences in approach, this study also departs from mainstream studies in more specific ways—reflecting, for the most part, an effort to deal with some of the methodological issues (mainly having to do with measurement) outlined in the prior section. I will introduce these elements below as they become relevant.

Data and Key Variables

The data for this study are drawn from school districts in the state of California. The current reality is that California and all other states administer their own achievement tests, which usually provide good measures of what students are learning—and better measures than can be obtained from most national data sets, which, if they have test scores at all, typically use a much smaller number of test items, focus only on students in certain grade levels, and have few students per district. The problem is that states tend to use different tests that cannot readily be compared to one another. So while an analysis that is national in scope might otherwise be desirable, there are grounds in this case for picking a particular state. And California offers a nice compromise. It has good measures of student achievement whose reliability has been carefully evaluated and maintained over time. It is also a large, diverse state whose districts vary greatly on variables of possible relevance—they are large and small, urban and rural, high minority and low minority, and so on—giving us a broad base for analysis.

Since 1998, when its school accountability plan went into effect, California has been administering achievement tests to its students and, based on the scores of all students across all grades in all subjects tested, giving each school a performance score called the Academic Performance Index (API). The test scores of individual students are confidential. But the API scores of schools are very public indeed, and are the state's prime means of holding schools accountable. They are also useful for the type of analysis we are carrying out here: they provide an overall index, for all students in each school, of how much students are learning. This is an attractive alternative to using the math scores of 4th graders or the SAT scores of college-bound high school seniors.⁴

⁴ Information on API scores can be found on the California Department of Education web site, at www.cde.ca.gov. For 1998-99 and 1999-2000, the API annual achievement scores were based entirely on the Stanford 9 achievement test, which covered reading, language, spelling, and math in grades 2-8, and reading, language, math, science, and social science in

The collective bargaining contracts were gathered randomly from 371 of California's approximately 1000 school districts. The contracts are filled with rules. Some of these rules deal with teacher pay, benefits, and time off, but most—the focus here—are designed to impose structure on the workplace by giving teachers various formal rights and restricting managerial control. To allow for a meaningful coding of the contracts, I singled out various work-rule dimensions as potentially important to the everyday operation of schools and typical of the kinds of restrictions unions fight for. Among them, for example, are rules regarding the assignment of teachers to classes, the voluntary and involuntary transfers of teachers to schools, and the evaluation of teachers. Factor analysis was employed to reduce all coded dimensions to a single index of restrictiveness. (See Appendix.)

The empirical analysis to follow, then, is built around the two key variables I have just introduced. One is California's index of school performance, the API, which is derived from student test scores. The other is an index of the restrictiveness of the collective bargaining contract, which is derived from a coding of contracts. The question at issue is: how does restrictiveness affect school performance?

The Empirical Model

The purpose of California's accountability system is to improve the performance of the public schools and thus to boost student achievement. This is the goal of public policy. It makes good sense, therefore, to take as our dependent variable the growth in API scores over time, and to frame the empirical analysis as one of determining whether the restrictiveness of collective bargaining contracts makes it more difficult for public schools to achieve increases in their API scores.

grades 9-11. In subsequent years, the department shifted gradually away from the Stanford 9 toward its own achievement tests in calculating the API annual scores.

The model of school improvement I'll be testing here consists of three basic parts. The first recognizes that how much a school improves is likely to depend on its initial level of performance, the base API. One reason is that schools starting out at very low achievement levels have vast opportunities for improvement—nowhere to go but up—while schools starting out at high achievement levels may have to struggle to increase their scores at all, not to mention by amounts comparable to the lower level schools. A related reason is that the API scale, which by design runs from 200 to 1000, gives rise to ceiling effects as scores begin to approach 1000. While only about 10 percent of the schools in this sample have scores over 800, schools at this level do not have the same scope for improvement as schools starting out at lower levels. The model includes a quadratic term (the square of the base API) to take this nonlinearity into account.

Controlling for a school's starting point, its academic improvement depends on characteristics of the students it is trying to educate, as well as characteristics of the school itself and the district it is a part of. These influences are best incorporated in two separate ways, which give rise to the second and third components of the model. The first has to do with the basic levels that these variables take on for particular schools. Schools with larger percentages of minority students, for example, may find it more difficult to improve achievement than schools that are mainly white. Similarly, schools that have larger enrollments, larger classes, or more inexperienced teachers may find it more difficult to improve than schools that are more advantaged on such counts. The second component recognizes that, whatever the levels of these variables for a given school, the school's change in performance over a given time period may also be a function of how these variables have changed over that period. If the student body has shifted from 10 percent to 20 percent Latino, or if enrollment has jumped from 500 to 600, these changes might be accounting for some of the movement in student achievement. The model to be estimated therefore takes the following form:

$$\text{APIGrowth}_{ij} = \beta_0 + \beta_A \text{API}_{ij} + \beta_{AS} \text{API}_{ij}^2 + \beta_L X_{ij} + \beta_C \Delta X_{ij} + \beta_R R_{ij} + \varepsilon_{ij}$$

The subscript i is a counter for schools, and j is a counter for districts. X is a vector representing the level that the (nonunion) independent variables take on during the base period. ΔX is a vector representing the change in these X variables over the relevant time period. R is the restrictiveness of the collective bargaining contract. And ε is the random error component. There are several measurement and estimation issues worth discussing here, so let me add a few points of clarification.

The base year of this analysis is the 1998-99 school year: the first year of the California accountability program and also the year in which the collective bargaining contracts were obtained from the districts. Achievement growth is measured as the total improvement in API scores between this base year and the 2002-03 school year. This five-year period of time is chosen because it is long enough to give the schools an opportunity to demonstrate improvement (or not) but short enough to help assure that the restrictiveness of the union contract, which is measured only in the base year, remains relatively stable.⁵

The growth in school performance is a measure of how each school's API scores have changed over the period. It is not, however, the simple difference between the API in 2002-03 and the API in 1998-99. Because of changes in the achievement tests and the rules for taking them (regarding, e.g., which students can be excluded), the Department of Education made adjustments in each year's API scores to assure that they could be validly compared to the scores of the immediately prior year to yield yearly growth scores. Unique adjustments were carried out each year. The total growth score for any given school over the entire period, then, is determined by summing its four year-to-year growth scores.

⁵ The term of most contracts is about three years, and, although many will not change much with renegotiation (or may take a year or two to renegotiate), the reliability of the restrictiveness measure is bound to decline the longer the time period. A period of five years seems a reasonable compromise under the circumstances.

The analysis is conducted at the school level, and it is carried out separately for elementary and secondary schools to recognize basic organizational differences between the two. In elementary schools, for instance, teachers usually have the same students all day long and have very personal relationships with them, whereas in secondary schools relationships are less personal and more bureaucratic, so it is reasonable to think that collective bargaining (and other variables as well) may operate somewhat differently in these settings.

Whichever type of school is being analyzed, we have to recognize that the schools are clustered into districts, and that, due to a host of unmeasured variables that make up their economic, social, cultural, and governing environments, schools within a given district have much more in common with one another than they do with other schools. The error terms in this analysis are therefore likely to be correlated within districts, and this violates the usual OLS assumptions. This being so, the analysis is carried out using a robust (Huber-White) estimator of variance that recognizes the within-district correlation of errors across schools (clustering).

Endogeneity bias is always a concern, but it seems unlikely to be a problem that demands correction here. There is some possibility that achievement has a causal effect on restrictiveness, e.g., with teachers pressing for more protections when students are performing at low levels. But our dependent variable is the change in achievement, not achievement per se, and this should lessen any problems of bias—particularly because the change in achievement occurs in the years after the labor contracts have been negotiated, and cannot have caused their restrictiveness. We are also explicitly controlling for the base year level of achievement on the right hand side, as well as for an array of student background factors. The correlation between achievement and restrictiveness, moreover, is just $-.02$ for elementary schools and $.06$ for secondary schools.⁶ Another possibility is that endogeneity

⁶ These correlations are at the district level. The level of contract restrictiveness in 1999 is due to the accumulated influences of factors in the past, so contemporaneous correlations can only be suggestive. Even so, it is worth noting that restrictiveness

problems may arise from factors that are causally connected to achievement and also correlated with restrictiveness but omitted from the model. We might worry, for example, that highly bureaucratic systems are more likely to have restrictive contracts, and that bureaucracy creates conditions that lead to bad teaching; or we might worry that unions are more powerful and get more restrictive contracts in large, urban, high-minority systems that tend to be poor performers. But our model includes a long list of student, school, and district controls, including measures of bureaucracy, teacher quality, district size, bureaucracy, and student ethnicity. We cannot know that all influences of any relevance are included, but our controls are quite extensive, and it is reasonable to think that any omitted factors have rather small effects.⁷

Independent Variables

As set out above, our model contains a number of unspecified X's that are potential determinants of school improvement. Let's take a more detailed look at these X variables.

Research on student achievement has arrived at one central conclusion that is essentially beyond dispute: the characteristics of the students—socioeconomic status, ethnicity, language problems—are the key determinants of achievement, and are consistently more powerful influences than characteristics of schools and districts (Hanushek, 2002, 2003). In our model, the following variables are designed to capture the composition of each school's study body:⁸

is contemporaneously correlated at much higher levels with other variables, e.g., district size (.62), rural location (-.40), and percent black (.31).

⁷ It is de rigeur among some researchers to "correct" for endogeneity problems if there is any hint they might exist. But corrections introduce new problems of their own that can easily lead to poorer rather than better estimates. Good instruments are often difficult to find, and that is the case here. The most obvious candidate is district size, which is strongly correlated with contract restrictiveness; but as the analysis later shows, it may also be connected in various ways to student achievement, which rules it out as an instrument. Other candidates are the percent of district voters who are Democrat, or the percent of district adults who are unionized (or belong to public sector unions); but these turn out to have very weak correlations to contract restrictiveness, making them weak instruments.

⁸ All of the data for these measures are obtained from the API data bases, which can be found (as of March, 2007) at www.cde.ca.gov/ta/ac/ap/apidatafiles.asp.

- Student ethnicity: percent black, percent asian, percent Latino, percent other minority (with percent white the omitted category).
- Student socioeconomic status: percent qualifying for free or reduced-price meals.
- Student language problems: percent categorized as English language learners.

Research on school and district characteristics, by contrast, has not led to a coherent set of findings. It might seem that higher district spending, smaller school size, smaller class size, higher teacher salaries, and other common-sense factors would have strong positive connections to how much students learn, but studies have not been able to document as much, and the findings have generally been quite mixed (Hanushek, 2002, 2003). There is recent experimental evidence that smaller classes may have a modest impact on student learning, at least in the early years (Mishel and Rothstein, 2002). There is also good evidence that teacher quality is the most important organizational determinant of how much students learn (Sanders and Rivers, 1996). Teacher quality, however, is difficult to measure in the absence of very detailed data. One proxy for (low) teacher quality is the percentage of inexperienced teachers in a school, for it has been shown that teachers with only one or two years of experience in the classroom are not as effective as more experienced teachers (Rivkin, Hanushek, and Kain, 2001). Another proxy is the percentage of teachers with credentials, although the evidence linking credentials with student performance is weak (Kane, Rockoff, and Staiger, 2006).

In the interests of having adequate controls, but without much expectation that most of the school and district variables will show even modest relationships with school improvement, I have included in the model a rather lengthy list of such variables for which measures are available. They are:⁹

⁹ The data on class size are taken from the API data sets, referenced earlier. All other data are from the California Department of Education's CBEDS data files, except for (1) the data on teacher salaries, total spending, and bureaucracy, which are from the Department's J-series financial data files, located on the web at www.cde.ca.gov/ds/fd, and (2) the data on the education level of the district population, which are from the National Center for Education Statistics' Census 2000 School District Demographics Data Files, which can be found on the web at nces.ed.gov/surveys/sdds/selectgeo.asp.

- District size: the log of district enrollment.
- School size: the log of school enrollment.
- District location: dummies for rural and suburban (with urban the omitted category).
- District education level: the percentage of adults living in the district who have graduated from college.
- Type of district: dummies for high school district and elementary district (with unified district the omitted category).
- Type of school: dummy for high school (in secondary school analysis).
- Class size: for elementary schools, average class size for grades K-3 and for grades 4-6 (two separate measures); and for middle and high schools, average class size of core academic classes.
- Teacher inexperience: percentage of a school's teachers who have been teaching for less than three years.
- Teacher credentials: percentage of a school's teachers who are fully credentialed.
- District spending: log of district total expenditure per student.
- Teacher salaries: log of teacher salary at "step 10" (an identifiable mid-level category).
- Administrative overhead: ratio of total spending on administrative positions to spending on teacher salaries.

Basic Findings

Before estimating the model, let's take a descriptive look at some of the key variables. The index of contract restrictiveness ranges from 0 and 6, with a mean (over the districts included in the analysis) of 3.23 and a standard deviation of .77. A common notion is that large districts tend to have much more restrictive labor contracts than small districts do; and the scatterplot in Figure 1A reveals that there is indeed a marked bivariate relationship between district size and the restrictiveness of the contract ($R^2=.40$). Administrators and authorities are confronted with much more restrictive formal contract rules in the larger districts, and teacher rights and autonomy are far more expansive.

One might think that labor contracts would also be especially restrictive in districts with high percentages of minority students, if only because, in the nation as a whole, large districts tend to be high-minority districts. But in California this is not the case. Most minorities in California are Latino, and they are well represented in small as well as large districts throughout the state. As Figure 1B shows, there is only a slight bivariate relationship between minority student enrollment in a district and the restrictiveness of the bargaining contract. Generally speaking, minority kids in California are no more likely to attend rule-bound districts than nonminority kids are.

Finally, let's take a look at the relationship that most concerns us here: the one between contract restrictiveness and the growth in student achievement. Because achievement growth has been higher in California's elementary schools than in its secondary schools, scatterplots are presented separately for each. Figure 2A shows that, for elementary schools, there is no bivariate relationship between restrictiveness and achievement at all. For secondary schools, on the other hand, Figure 2B shows a bivariate relationship consistent with our central hypothesis: contract restrictiveness is associated with lower achievement growth.

It is always informative to look at the raw data to get the initial lay of the land. But simple two-way relationships of this sort may obviously be misleading, because they fail to take into account the effects of other variables. In the presence of controls, the relationship between restrictiveness and achievement growth may turn out to be quite different than it appears in these figures.

The full model is estimated in Table 1. Impact scores are computed to provide a common metric for measuring the effect of each variable on school improvement. These impact scores answer the question: by how many standard deviations does school improvement change when the relevant independent variable shifts from a low value (one standard deviation below its mean) to a high value (one standard deviation above)?

The results make it clear, as research has repeatedly shown, that the most powerful impacts on school performance are the background characteristics of students. Controlling for their initial levels of achievement, both elementary and secondary schools are more likely to improve their performance the lower the percentage of African-American and Latino students, the higher the percentage of Asian students, the lower the percentage of students on free and reduced-price meals (the measure of SES), and the lower the percentage of students who have difficulty with English. As measured by their impact scores, the effects of these variables on achievement growth are considerable, roughly in the range of .30 to .40 standard deviations of improvement. The impacts would be even higher, moreover, were the background variables not competing with one another for influence, for they are partially measuring the same things. A school with a high percentage of Hispanic students, for example, also tends to have high percentages of kids on free and reduced-price lunch and kids with language problems. Were one or more of these variables omitted from the equation, the estimated impact(s) of the remaining variable(s) would be even larger, and the few with weak or near-zero impacts would rebound substantially.

As prior research also leads us to expect, the school and district variables are anemic by comparison. School size and the percentage of inexperienced teachers have statistically significant impacts at the elementary level. But their impacts, while in the right direction, are not significant for secondary schools; and for elementary schools, the change component of the size variable has a significant impact in the wrong direction. Most of the other school and district variables that might plausibly be related to achievement gains—district size, district spending, teacher pay, bureaucracy—do not have significant effects at either level.¹⁰

Two “significant” findings are exceptions, but probably spurious. Specifically, for both types of schools, achievement increases when class size goes up. Also, for elementary schools, achievement increases when district spending goes down. There is neither a research base nor a theoretical rationale to support either of these results. The explanation may simply be that, in equations with many interrelated variables, there is enough multicollinearity to produce odd results from time to time.

Overall, the school and district variables are a disappointment, as expected. In light of this result, and in light of the literature’s consistent inability to document reasonable connections between the organization of schools and the achievement of students, it is notable that the restrictiveness of the collective bargaining contract does indeed seem to have such a connection. As Table 1 indicates, the restrictiveness of the contract has a statistically significant impact on achievement growth in both elementary and secondary schools, estimated separately. The direction of the impact is precisely what we ought to expect: it is negative, making achievement gains more difficult. And its impact scores, -.24 for elementary schools and -.32 for secondary schools, are greater than those of any other organizational variables in the analysis.

¹⁰ Chubb and Moe (1990) found that bureaucracy does have a significant impact on achievement, but their measure of bureaucracy was based on measures of administrative influence, whereas the measure employed here is based on spending. (Influence items are not available in the California data.)

One complication, which is inevitable, is that contract restrictiveness presumably affects student achievement by affecting the organization of schools and districts, represented here by such variables as class size, inexperienced teachers, credentialed teachers, and bureaucracy, which we are controlling for as potential determinants of achievement. By including these variables in the analysis, however, some of the impact of restrictiveness may be channeled through these factors, and not picked up in its own coefficient, which represents only its direct effect on achievement. It would take a much more extensive analysis to sort all this out—constructing models of class size, models of teacher inexperience, and so on—and I will leave such work to future research. For present purposes, I will simply point out that these other variables have precious little impact in this analysis on the outcome variable anyway. And when models are estimated in which these variables are excluded entirely, the impacts of contract restrictiveness on achievement remain virtually unchanged, as do their significance levels.¹¹

Conditional Effects

So far, the estimation indicates that contract restrictiveness has a negative effect on student learning in both elementary and secondary schools, and that it is more influential than any other organizational variable in the model. The model we have employed, however, assumes that collective bargaining has an effect on student achievement that is the same under all conditions. This is the most straightforward assumption to make, and it is the reasonable place to begin. But reality may be more complicated, and the impact of collective bargaining may actually vary depending on the conditions in which it operates.

¹¹ Ideally, this would not be the case. It may occur here because the full array of important organizational variables is simply not represented or well enough measured. It may also be because the relationship between contract restrictiveness and some of these variables is conditional and/or nonlinear—which is surely the case for inexperienced and credentialed teachers, because transfer rules (a basic component of the restrictiveness index) may lead to more experienced and credentialed teachers in some schools (affluent ones) and less experienced and credentialed teachers in other schools (disadvantaged ones).

In the theoretical section, I argued that the consequences of collective bargaining for achievement should tend to be negative, at least on balance. This argument is rooted in fundamentals—in the basic interests of unions and teachers and the kinds of restrictions they are led to pursue—and its expectations are still the core expectations here. But if we move beyond fundamentals and consider specific conditions that might affect how collective bargaining actually operates, these core expectations could be modified, perhaps significantly. Collective bargaining could prove to be quite negative for achievement under some conditions, but not so negative under others, and even positive under still others—pointing the way toward a more finely grained understanding of union impact.

Because an exploration of conditional effects takes us into new territory that could easily get quite complicated, and because space is limited, I will focus here on just two aspects of public schooling that seem good candidates for qualifying the impacts of collective bargaining. One is the size of the school district. The other is the minority composition of the school.

In California, as in most states, districts of all sizes are unionized and engage in collective bargaining. But it is plausible to suspect that formal rules, including union contract rules, may in practice be less binding and thus less consequential in smaller districts. For in smaller districts, the people involved in personnel decisions—district leaders, administrators, teachers—are more likely to know one another and function as a community. This being so, they may sometimes agree to ignore or circumvent formal rules when their enforcement would conflict with what is good for children. As districts get larger, and especially as they get very large, relationships among participants should tend to get more impersonal and rule-governed, and formal contract rules would tend to be followed even if the consequences for children are clearly not good. This argument may or may not be correct, but it is well worth considering, and it leads to a hypothesis about conditional effects: that the consequences of collective bargaining will vary with district size, and will be more negative the larger the district.

Now consider the minority composition of the school, which may be relevant for several reasons. A rationale that has gained attention in the research literature is that rule-based standardization by unions may be good for “average” children, but ill-suited to children from disadvantaged backgrounds, who need more specialized attention (e.g., Eberts and Stone, 1984, 1986; Milkman, 1997). I am not persuaded by this argument, because union rules are pegged to the interests of teachers, not the needs of the average child. They easily come into conflict with the education of all children, not just those who are disadvantaged.

A more persuasive rationale is that the parents of children in high minority schools are likely to be less educated, less affluent, and less politically active than parents of children in other schools, and thus are likely to be less influential—with district officials, with principals—in preventing the implementation of contract rules that affect their own schools adversely.¹² By this logic, high minority schools may be more negatively affected by contract rules than low minority schools, even if the formal levels of restrictiveness they face are exactly equal. These effects may show up across districts, with the highest minority schools in the state being the most negatively affected (for a given level of restrictiveness). They can also show up within districts, where what counts is a school’s relative level of minority composition compared to other schools in its same district.

Minority composition is also relevant because of its connection to how teachers get distributed across schools (within districts). Studies have shown that teachers tend to leave high minority for low minority schools when they can, making it more difficult for disadvantaged schools to attract and retain the kinds of experienced, quality teachers that all schools are looking for (Hanushek, Kain, and Rivkin, 2004; Levin, Mulhern, and Schunck, 2005). Seniority based transfer rights, which are common in the

¹² Note this is about the implementation of union rules, not their adoption. Collective bargaining occurs behind closed doors, and the rules contained in contracts are probably unknown to virtually all parents. Specific rules only become relevant to parents when they result in unpopular decisions that affect their own schools, and it is at this stage—the implementation stage—that one would expect differences in parent influence to have an effect.

more restrictive union contracts, are likely to exacerbate this problem: giving teachers formal rights to transfer from less desirable to more desirable jobs, and limiting the capacity of administrators to put their best teachers in schools with the greatest needs. Thus, within districts, restrictive contracts may have more negative effects on those schools that have relatively high levels of minority enrollment.

Finally, we need to recognize that, if the rationale behind the size hypothesis is valid—and thus if contract rules tend to be more binding in the larger districts—then the conditioning effects of minority composition are themselves likely to vary with district size. It would be in the larger districts that the weak influence of minority parents would tend to burden high minority schools with more onerous rules, and in the larger districts that transfer rights would have greater effect. If the logic is correct, district size has an overarching role to play.

Although tests for the conditioning effects of both district size and minority composition can be carried out within the same model, it is useful to begin the analysis by focusing first on district size—because in a few quick steps, we are led to results that are quite remarkable and instructive. Step 1 involves a simple reestimation of the basic model employed in the prior section, augmented with a new term that interacts restrictiveness of the labor contract with district size. The findings, set out in top portion of Table 2 (which, for simplicity, only presents the results for the key variables of interest here), shows that the interaction term is negative and significant for both the elementary and secondary school samples. It would appear from these findings that the effects of collective bargaining do indeed vary with district size and, as expected, become more negative the larger the district.¹³

But now suppose we separate out the larger districts from all the other districts, where the former are those enrolling more than 20,000 students. Reestimation should show the impact of restrictiveness to be more negative for the larger districts than for the others. And this is what we find—see Table 2—

¹³ Note that, while the coefficient of the restrictiveness variable itself is positive in both equations, the total impact on school improvement of a change in restrictiveness is always negative, due to the (negative) contribution of the interaction term.

but with the surprising twist that, for the smaller districts, the impact of restrictiveness is essentially zero. This is the case, moreover, in both the elementary and secondary samples, estimated separately.¹⁴

These results are incomplete, of course, but they do suggest that contract restrictiveness may not make much difference for student achievement except in the larger school districts. The notion that contract rules are not entirely binding, and can be gotten around when they conflict with the needs of students, thus finds empirical support. This is all the more impressive given that, in the state of California, districts with fewer than 20,000 students make up some 94 percent of all districts.

But the flip side is impressive as well. While districts with more than 20,000 students are much fewer in number, they actually enroll a whopping 47% of all students in the state. And in these districts, the impact of contract restrictiveness is much more negative than our earlier estimates implied. A shift from low restrictiveness to high restrictiveness leads to a drop in achievement of .44 standard deviations among elementary schools and .57 standard deviations among secondary schools. (The impacts were .24 and .32, respectively, in the prior analysis.) These effects are comparable in magnitude to those associated with student background characteristics, and tower over those associated with other aspects of school and district organization.

Given these findings, it makes sense to proceed with the analysis by looking separately at the larger and smaller districts. In each case, the base model is augmented by interaction terms that test for whether the effects of contract restrictiveness are conditional upon the minority composition of the school and the size of the district.

(1) District size is included to allow for the possibility that, even within size categories, size may have a more finely grained conditioning effect on contract restrictiveness that still needs to be taken into

¹⁴ While we are limited here by the relatively small number of districts in the sample with more than 20,000 students (39) and between 10,000 and 20,000 students (39), regressions using different cut-off points suggest that 20,000 is about the level at which collective bargaining proves to make a difference.

account. This is especially plausible within the set of larger districts, because that category is unbounded at the higher end and contains a greater range of sizes.

(2) Minority composition is defined by reference to the percentages of African-American and Latino students in the school. As these are separately controlled for in the basic model, one way to proceed here is to include separate interaction terms for each minority group. A simpler approach is to combine them into a summary interaction term, using percent minority. I estimate models for both.

(3) The above terms measure the “absolute” minority composition of the school. The “relative” minority composition of the school is measured as the difference between a school’s percent minority and the percent minority of the median school in the district.¹⁵

I should note that “relative” minority composition, whose postulated effects derive (in effect) from the intra-district competition among schools for resources, is at something of a disadvantage here. By far the best test of its impact is likely to be found in the elementary-level analysis of large districts, because its median district has 26.5 elementary schools—plenty for competition. In the secondary-level analysis, by contrast, even the larger districts have very few schools—a median of six middle schools and four high schools—and much of the competition is circumscribed: teachers typically don’t transfer from middle schools to high schools, or vice versa. For the smaller districts, the situation is much worse. In the secondary analysis, there is a median of just one middle school and one high school per district, which rules out most competition; and in the elementary-level analysis, the smaller districts have a median of just four elementary schools—better, but not by a lot. Although I will include the “relative” minority composition variable in all these analyses, it is in the elementary analysis of larger districts that the test is most meaningful.

Now let’s turn to the findings, set out in Table 3 (again, for simplicity the full set of control variables is not presented). Consider first the results for elementary schools in the larger districts.¹⁶ The

¹⁵ The main effect here is included as a control variable, as it was not in the original equation.

key interaction terms behave as expected. In Model 1, we find that the impact of contract restrictiveness on student achievement is more negative the higher a school's percentage of African-Americans, the higher its percentage of Latinos, the higher its minority composition relative to other schools in the district, and the larger the district is. All these conditioning impacts are statistically significant. In Model 2, which uses percent minority rather than the two separate minority measures, the conditioning effect of a school's "absolute" minority composition remains negative and statistically significant, as does district size; but the effect of the "relative" measure of minority composition, while still negative, drops a bit and just loses its statistical significance.

The findings are quite different for elementary schools in smaller districts. In Model 1, the effects of restrictiveness are significantly more negative for African-American students, but this is the only significant interaction in either Model 1 or Model 2. And we need to remember that the vast majority of minority kids in California are Latino, not African-American. If we look at the impact of percent minority in Model 2, it is negative, but it is statistically insignificant and much smaller in magnitude than the corresponding negative impact in the larger districts. The coefficient of the "relative" minority composition variable (at a disadvantage here) is positive and insignificant, as is the coefficient for district size. Thus, it is in the larger districts that we find conditioning effects for minority composition (both absolute and relative) that are in line with what we would expect, while in the smaller districts the results tend to suggest that not much of any consequence is happening, except for African-American students.

In the analysis of secondary schools, the results are very similar for the smaller districts. The coefficient for African-American students is big and negative; and although it is statistically insignificant, it raises a red flag—viewed together with the elementary results—that perhaps union rules

¹⁶ As in the prior table, the positive coefficient on the restrictiveness variable itself is does not mean that the overall impact of a change in restrictiveness is positive, for the overall impact also depends on the (negative) contributions of the interaction terms. See Table 4 for the calculation of impacts.

do work against the achievement of black students even in the smaller districts. In other respects, however, restrictiveness shows no impacts. All the variables are statistically insignificant. For the larger districts, on the other hand, minority composition clearly does have an impact. In Model 1, the coefficient for Latinos is negative and statistically significant—an important result, given the prevalence of Latinos in California schools. And although the coefficient for African-Americans does not achieve significance (in part a reflection of their low numbers), it is negative too and of almost the same size. When minority composition is simply measured as percent minority, as in Model 2, its coefficient is negative and highly significant. “Relative” minority composition is insignificant in both models, but that is perhaps to be expected, given the fairly small numbers of secondary schools within each district.

Table 4 illustrates how the impact of collective bargaining varies with district size and minority composition, based on the estimates in Model 1 for the larger districts.¹⁷ The most benign effect of collective bargaining arises when schools are in districts at the lower end of the size distribution, low in minority enrollment, and advantaged relative to other schools in their districts. With this combination, the impact of collective bargaining is close to zero (at .03). As conditions change—as districts get larger, as schools enroll more minorities—the effect of collective bargaining gets progressively more negative, until it reaches a substantial -.61 for schools that are in the larger districts, high in minority composition, and disadvantaged relative to other schools in their districts. A perusal of the table shows that this total change in impact (from .03 to -.61) can be broken down in the following way: -.21 of it is due to the increase in district size, -.32 is due to the increase in minority composition, and -.11 is due to the increase in intradistrict disadvantage. We have to remember, however, that these calculations are all based on schools in the larger districts, and that the impact of collective bargaining for schools in the

¹⁷ At the elementary level, calculations are carried out for eight different sets of conditions: depending on whether the district is near the lower or higher end of the size distribution, whether the school is low minority or high minority in composition, and whether it is advantaged or disadvantaged relative to other schools in the district based on its minority enrollment. For each variable, the lower cutoff point is the 25th percentile of the larger-district sample, the upper cutoff point is the 75th percentile of the larger-district sample.

smaller districts is essentially zero. The conditioning effect of district size, therefore, is much greater than the numbers in Table 4 suggest.

For secondary schools, the calculations are simpler, because the earlier estimation showed that district size does not have additional, more finely grained impacts within the set of larger districts, and that the “relative” minority composition of the school does not matter either. Thus, we need only carry out the calculations for changes in minority composition.¹⁸ When this is done, as Table 4 indicates, we find that the impact of collective bargaining on student achievement is a substantial -.40 even for schools with low minority enrollment, and that it jumps to a much more negative -.84 for schools with high minority enrollment. In general, not only does minority enrollment make a big difference here, but it appears that the negative effects of collective bargaining are much greater for secondary schools than for elementary schools. A high minority secondary school is more seriously affected by restrictive labor contracts than a high minority elementary school.

In sum, then, the empirical results indicate that the impact of collective bargaining on student achievement gains is not constant across schools, but varies depending on the size of the district and the minority composition of the school. This is true for both elementary and secondary schools, estimated separately. In each case, the restrictiveness of the contract does not appear to have much effect in the vast majority of school districts, but in the larger districts—which enroll roughly half of the state’s students—it has a very negative impact indeed, especially at the secondary level, and the magnitude of this impact is more pronounced for high minority schools.

¹⁸ This is carried out using Model 1 from Table 3, assuming that district size is at its median (within the set of larger districts) and that the “relative” minority composition variable is equal to zero (and thus at the median in minority composition for the school’s district).

Conclusion

This paper is about public education, but it illustrates a general phenomenon that is relevant throughout much of American government: that ordinary public employees, by getting organized for collective action, can exercise power in pursuit of their own special interests—and in so doing, can have important impacts on the policies, structures, and performance of government.

In general, these impacts are brought about through two avenues of influence. First, public sector unions participate actively in politics, taking advantage of massive memberships, deep financial pockets, and extensive political organization to bring their interests to bear in electoral and policymaking arenas at all levels of government. And second—the subject of study here—these same unions represent their members in collective bargaining, and in that realm too they use their power to shape the structure and ultimately the performance of government.

Although public sector unions have been actively engaged in both avenues of influence for more than a quarter century now, students of American government have paid little attention to them. The literature on bureaucracy recognizes that ordinary bureaucrats can gain power by having expertise unavailable to superiors, and that bureaucratic leaders can gain power by taking entrepreneurial action of various sorts. But scholars have little to say about the capacity of ordinary bureaucrats to exercise power through collective action. Similarly, the literature on interest groups focuses on a vast array of groups that arise outside of government to promote the interests of myriad social, economic, and political constituencies, but not on groups that arise from inside government to pressure on behalf of public employees. Groups like the National Education Association (NEA), the American Federation of State, County, and Municipal Employees (AFSCME), and the Service Employees International Union (SEIU)—which have millions of members and are major forces in American politics as well as in collective bargaining—are barely on political scientists' radar screens.

This is an oversight with important consequences. Some of the consequences are intellectual: it is the job of political scientists to understand the structure and performance of government, and they cannot do this adequately if they fail to recognize and explore the collective power of the government's own employees. Some of the consequences are more substantive: attempts by policymakers to build effective government agencies, or to improve those that are not effective, must ultimately rest on an understanding of what the determinants of agency performance are—and this is not possible if public sector unions are left out of the equation, as though they are somehow irrelevant to how government is organized and does its work.

Public education is a case in point. For more than two decades, improving the academic performance of the public schools has consistently been a top priority of American government. Throughout, reformist attention has especially centered on large, urban school districts, which typically have very high concentrations of minority students and are often abysmally ineffective at educating them. There are doubtless many factors that contribute to these problems. But it is only reasonable to ask whether public sector unions, by imposing structures that are literally not designed to promote the academic achievement of students, are at least part of the explanation for why the schools are not doing their jobs very effectively and why they are so difficult to improve. There can be no answer if the question is not even raised, and if political scientists and policymakers continue to think about government in ways that have nothing to do with collective power of public employees.

The analysis of this paper is not definitive, of course, but it moves the ball downfield. First, it suggests that teachers unions do matter for the performance of the public schools. In particular, it supports (with qualification) the core expectation that, because union and teacher interests are often in conflict with the educational interests of children, the restrictions built into labor contracts should on balance tend to have negative consequences for academic achievement. Second, in actively exploring

the possibility that the effects of collective bargaining may depend on the conditions in which it operates, the analysis offers a more finely grained understanding of the connection between collective bargaining and the schools.

- Collective bargaining appears to have a strongly negative impact in the larger districts, but it appears to have no effect in smaller districts (except possibly for African-American students—which is important indeed if true). This supports the conjecture that, because smaller districts are more like communities and larger districts more bureaucratized, formal contract rules are more likely to be enforced—and to have consequences—in the larger districts.
- Among the larger districts, the negative effects of collective bargaining are much greater in magnitude for high minority schools than for other schools. This is consistent with the argument that minority parents are less politically influential, and less able to step in when enforcement of formal rules would affect their schools perversely.
- Although the findings are more mixed on this count, the best evidence indicates that the impact of collective bargaining is especially negative for schools that are “relatively” high minority within a given (larger) district. This supports the argument that restrictive contracts put high minority schools at a disadvantage in the competition for teachers and resources within districts.

If these findings are essentially on the mark, their substantive implications are quite important.

For what they tell us is that collective bargaining does have negative consequences for student achievement, and that the effects are concentrated on precisely those districts and schools—large districts, high-minority schools—that, over the years, have been the lowest performers and the most difficult to improve. It follows that efforts to boost achievement in these disadvantaged contexts, as well as to reduce the longstanding achievement gap between whites and minorities, need to recognize that collective bargaining may well be part of the problem—and that it deserves to be taken seriously as a

target of reform. There is no magic bullet here. But the evidence suggests that systematic attempts to reduce the restrictiveness of labor contracts could have significant payoffs for public education.

More research is needed, of course, to be confident about these findings and implications. That is always the case for any empirical project, but especially one on a subject that, like this, is so little studied. The greater need, however, is for broadly based research on the power of public sector workers generally—in collective bargaining, in the policy process, in politics generally—and for theoretical perspectives that link their power to an understanding of government.

Appendix

The collective bargaining contracts were coded by singling out work-rule dimensions of relevance to school performance. Some were measured by single indicators, some by multiple indicators. The dimensions and their indicators are listed below. The numbers in parentheses are their factor loadings on the relevant dimensions.

- Assignment of teachers to classrooms (.30)
 - Role of seniority in assignments. (.70)
 - Can teacher formally request reasons for assignment? (.70)
 - Do teacher preferences have to be formally considered? (.40)
 - Stipulation that changes to assignments not be arbitrary and capricious. (.41)
 - Does teacher have a right to continue in current position? (dropped)
 - Are there time deadlines by which assignments must be made? (.58)
 - Number of pages devoted to assignments, logged. (.82)
- Voluntary transfers across schools (.68)
 - Role of seniority in voluntary transfers. (.62)
 - Are rejected candidates formally entitled to reasons for outcome? (.66)
 - Are there time deadlines for posting jobs and making job decisions? (.53)
 - Number of pages devoted to voluntary transfers, logged. (.62)
- Involuntary transfers across schools, as well as surplusings (school closings, etc.) (.78)
 - Role of seniority in involuntary transfers. (.63)
 - Can teachers formally request reasons for involuntary transfers? (.52)
 - Is there a limit on the number of involuntary transfers? (.40)
 - Are there time deadlines for involuntary transfers? (.39)
 - Role of seniority in the surplusings of teachers. (.68)
 - Are teacher preferences solicited in surplusings? (.48)
 - Do surplusd teachers have a right to return or priority in rehiring? (.47)
 - Number of pages devoted to involuntary transfers and surplusings, logged. (.74)
- Policy Committee with teacher representation? (.26)
- Personnel Committees with teacher representation (.25)
 - Is there a committee with input on assignments? (.47)
 - Is there a committee with input on voluntary transfers? (.52)
 - Is there a committee with input on involuntary transfers? (.48)
 - Is there a committee with input on surplusings? (.51)
- Observation of teacher performance (.24)
 - Is advanced notice required for classroom visits and observations? (.47)
 - Is a minimum number of classroom visits required? (.67)
 - Are certain minutes of observation required for classroom visits? (.62)
- Joint determination of goals by teacher and principal? (.31)
- Evaluation of teachers (.44)
 - Does teacher get some choice of evaluators? (.44)
 - Is a third party involved in cases of disagreement? (.44)
 - Number of pages devoted to teacher evaluation. (.46)
- Restrictions on number of students per teacher (.56)
 - Are there limits on class size? (.69)
 - Are there other limits on teacher-student ratios? (.69)
- Faculty meetings (.25)
 - Is there a limit on the number of faculty meetings? (.56)
 - Is there a limit on the length of faculty meetings? (.60)
 - Is there a limit on the number of meetings with parents? (.58)
 - Are there limits on the number of length of other kinds of meetings? (.44)

- Does suspension or dismissal require “good cause” or “just cause”? (.25)
- Is a certain amount of preparation time guaranteed? (.41)
- Level of reliance on arbitration. (.25)
- Are there limits on nonteaching tasks? (dropped)
- Parent complaints (dropped)
 - Are there procedures for dealing with parent complaints? (.81)
 - Number of pages devoted to parent complaints, logged. (.81)

The most obvious method for creating an index of contract restrictiveness is through summation. For each dimension (if measured by multiple indicators), the component items can be standardized and averaged to give a composite score for that dimension; and these dimension scores can then be standardized and averaged to yield an overall index for the restrictiveness of the contract. But this approach assumes that the components within a dimension are all equally important (and should thus be given equal weight) in measuring that dimension, and also that each of the dimensions is equally important (and should be equally weighted) in constructing the final index—which are probably not good assumptions. Some rules and some dimensions are probably more important than others in capturing the restrictiveness of the contract. Rules about teacher transfer rights, for example, would seem much more problematic for school managers—and much more restrictive—than requirements about the length of faculty meetings.

This being so, factor analysis is a better, more flexible way of constructing the index. For each dimension, we can assume that the component rules are all measures of some latent concept (such as the restrictiveness associated with voluntary transfer rights), and we can factor analyze them to create a single score for that dimension: a score that weights the components unequally, depending on their correlation with the latent variable. We can then assume that the fifteen dimension scores are all measuring (in different ways, some better than others) the restrictiveness of the overall contract. And by factor analyzing these fifteen scores, we can arrive at a summary index of restrictiveness that is based on all of them, with weightings reflecting their differential importance.

This was the approach adopted here, using the principal factor method within Stata 9.0. For each dimension measured by multiple items, factor analysis of the component items showed that one latent factor was overwhelmingly dominant. In the above list, the number in parentheses following each item is its loading on the common factor for its dimension. These loadings, as the numbers indicate, were generally in the .5 to .6 range. One of the items listed above—under the Assignment dimension, measuring whether teachers have rights to continue in their current position—was dropped because its loading was quite low (at .08). All the other loadings were at least .40, regardless of the dimension.

Based on these loadings, composite indexes (single scores) were then created for each dimension, and the fifteen dimensions were then treated as providing multiple indicators of the restrictiveness of the labor contract. Here again, factor analysis led to the emergence of just one underlying factor of consequence. (Its eigen value was 2.31, for example, while the second factor had an eigen value of just .38). The factor loadings across dimensions were much more varied and generally weaker in magnitude by comparison to the within-dimension loadings; but this is not surprising given the broad diversity of the measures employed here—an attempt to cover most all the different types of work rules—and the narrowness of some of them. Dimensions were eliminated if they did not have a least a .20 loading with the underlying factor; and two of the dimensions—the ones dealing with parent complaints and with nonteaching duties—were dropped from the analysis on this basis. Of the remaining thirteen dimensions, all have consequential roles to play in the overall index, but the most important are voluntary and involuntary transfers, followed by restrictions on the number of students per teacher, rules for teacher evaluations, and guaranteed preparation time. That transfer restrictions are more heavily weighted than the other work rules is perhaps an indication of just how constraining it is for the schools when they are bound by rules that make it difficult to choose their own teaching staff.

Based on these loadings, a single composite index was created to measure the restrictiveness of the labor contract. A constant of 3 was then added to this initial index to ensure (for simplicity of interpretation) that the

final index is always positive. This final index varies from 0 to 6, with a mean of 3.23 and a standard deviation of .77. This is the key independent variable in the empirical analysis of the paper.







Table 1: The Impact of Contract Restrictiveness on Achievement Growth

	Elementary Schools			Secondary Schools		
Variables	Coef.	(sd)	Impact	Coef.	(sd)	Impact
API, base year	-.31**	.075	-1.73	-.22**	.11	-1.48
API Squared	-.00*	.00	-.67	-.00	.00	-.38
District Size (log)	1.33	1.98	.07	-2.59	2.17	-18.16
Δ District Size (log)	2.81	18.00	.01	31.67	21.07	.16
School Size (log)	-11.05***	3.40	-.16	-6.16	3.89	-.20
Δ School Size (log)	10.71*	5.47	.07	8.90	6.54	.09
% Black	-.56***	.12	-.32	-.80***	.16	-.52
% Asian	.43***	.08	.23	.29**	.12	.20
% Hispanic	-.04	.10	-.04	-.21*	.12	-.30
% Other Nonwhite	-.05	.17	.01	-.07	.14	-.03
Δ % Black	-2.62***	.41	-.39	-2.45***	.45	-.38
Δ % Asian	1.24***	.34	.20	.62*	.35	.12
Δ % Hispanic	-1.06***	.24	-.25	-.93***	.32	-.25
Δ % Other Nonwhite	-.55**	.24	-.08	-.49*	.27	-.12
% Free Meals	-.37***	.08	-.46	-.24*	.13	-.35
Δ % Free Meals	-.04	.04	-.04	.01	.14	.00
% English Learners	-.46***	.12	-.41	-.09	.18	-.08
Δ % English Learners	-1.31***	.16	-.36	-.87	.26	-.30
Rural	-11.29	7.16	-.24	-8.52	6.57	-.23
Suburban	-1.78	2.89	-.04	-4.02	3.64	-.11
% College Educated	43.48***	13.70	.23	25.17*	15.16	.19
Elementary District	-6.24	3.82	-.13	----	----	----
High School District	----	----	----	-1.01	4.90	.03
High School	----	----	----	-20.78***	3.61	-.58
Class Size, Grades K-3	1.92**	.93	.13	----	----	----
Δ Class Size, Grades K-3	2.13**	.84	.16	----	----	----
Class Size, Grades 4-6	.89**	.43	.11	----	----	----
Δ Class Size, Grades 4-6	.17	.30	.03	----	----	----
Class Size, Core Acad.	----	----	----	1.25*	.66	.21
Δ Class Size, Core Acad.	----	----	----	.70	.59	.11
% Inexper. Teachers	-.25*	.15	-.11	.18	.25	.08
Δ % Inexper. Teachers	-.48***	.13	-.20	-.20	.19	-.10
% Credentialed Teachers	.14	.21	.06	-.26	.29	-.13
Δ % Credentialed Teachers	.24	.16	.09	.07	.18	.03
District Spending (log)	-42.76*	22.00	-.14	-3.07	21.03	.01
Δ District Spending (log)	-13.42	24.38	-.03	17.81	20.84	.07
Teacher Salary (log)	19.80	19.22	.07	23.58	18.78	.11
Δ Teacher Salary (log)	-39.61	29.41	-.11	2.53	24.29	.01
Bureaucracy	-26.23	52.70	-.03	-30.46	48.21	-.06
Δ Bureaucracy	-33.32	54.16	-.03	17.14	44.53	.03
Contract Restrictiveness	-7.56***	2.55	-.24	-7.65***	3.260	-.32
Constant	548.02***	261.72		99.87	259.24	
N Schools	1947			829		
N Districts	241			250		
Regression Adj. R ²	.56			.41		

Dependent variable is the growth in the API score from 1998-99 to 2002-03. Statistical significance is indicated as follows: *** p<.01, ** p<.05, * p<.10. Analysis carried out in Stata with clustering on the school district. All tests are two-tailed except for the test on contract restrictiveness, as we are testing a one-sided hypothesis in that case. Standard errors are in parentheses. Except for dummy variables, "Impact" refers to the effect on API growth, in standard deviations, of a shift in the relevant independent variable from a low value (one standard deviation below its mean) to a high value (one standard deviation above). For dummy variables, "Impact" has the same meaning, except it captures the effect of a shift in the independent variable from 0 to 1

**Table 2: The Impact of Contract Restrictiveness on
Achievement Growth, By District Size**

Elementary Schools					Secondary Schools			
Districts	Variables	Coef.	(sd)	Impact		Coef.	(sd)	Impact
All								
	Contract restrictiveness	11.81	11.32			23.44	14.17	
	Restrictiveness * district size	-2.01**	1.22			-3.31***	1.64	
	N Schools	1947				829		
	N Districts	241				250		
	Adj. R ²	.56				.44		
Larger								
	Contract restrictiveness	-16.31***	2.20	-.44		-15.65***	3.97	-.57
	N Schools	980				363		
	N Districts	32				33		
	Adj. R ²	.61				.54		
Smaller								
	Contract restrictiveness	-.91	2.63	-.02		-.41	2.63	-.02
	N Schools	967				466		
	N Districts	209				217		
	Adj. R ²	.57				.43		

Dependent variable is the growth in the API score from 1998-99 to 2002-03. Estimated model includes all the independent variables listed in Table 1. Statistical significance is indicated as follows: *** p<.01, ** p<.05, * p<.10. Analysis carried out in Stata with clustering on the school district. The test for contract restrictiveness in the first regression is two tailed, because the presence of the interaction term means that it could either be positive or negative and still be consistent with the hypothesis of an overall negative effect. The tests for all other variables are one tailed, as the hypothesis in each case is one-sided: that the coefficient is negative. Standard errors are in parentheses. "Impact" refers to the effect on API growth, in standard deviations, of a shift in the relevant independent variable from a low value (one standard deviation below its mean) to a high value (one standard deviation above). Impacts are not calculated for the first regression because the impact of restrictiveness on API growth depends on district size, and thus varies.

Table 3: Testing for Conditional Effects – Does the Impact of Contract Restrictiveness Depend on Minority Composition and District Size?

Dist. Size	Variables	Elementary Schools		Secondary Schools	
		Model 1	Model 2	Model 1	Model 2
		Coef.	(sd)	Coef.	(sd)
Larger					
	Restrictiveness	108.88**	42.96	117.34***	39.12
	Restrictiveness*black	-.32**	.18	-.18	.19
	Restrictiveness*Hispanic	-.14*	.10	-.25**	.11
	Restrictiveness*minority	----	----	-.19***	.08
	Restrictiveness*relative minority	-.18*	.11	.16	.16
	Restrictiveness*log district size	-10.64***	4.04	-11.45***	3.66
	N Schools	980		980	
	N Districts	32		33	
	Adj. R ²	.63		.54	
Smaller					
	Restrictiveness	-22.35	17.35	-16.84	17.14
	Restrictiveness*black	-.65*	.44	----	----
	Restrictiveness*Hispanic	-.05	.08	----	----
	Restrictiveness*minority	.19	.18	-.06	.08
	Restrictiveness*relative minority	----	----	.18	.19
	Restrictiveness*log district size	3.00	2.15	2.12	2.10
	N Schools	967		967	
	N Districts	209		217	
	Adj. R ²	.58		.43	

Dependent variable is the growth in the API score from 1998-99 to 2002-03. Estimated model includes all the independent variables listed in Table 1. Statistical significance is indicated as follows: *** p<.01, ** p<.05, * p<.10. Analysis carried out in Stata with clustering on the school district. The test for the contract restrictiveness variable itself is two tailed, because the presence of the interaction term means that it could either be positive or negative and still be consistent with the hypothesis of an overall negative effect. The tests for all the interaction terms are one tailed, as the hypothesis in each case is one-sided: that the coefficient is negative. Standard errors are in parentheses.

Table 4: Impact Coefficients for the Larger Districts

Type of School	Minority	Disadvantage	District Size	Impact
Elementary	High	High	High	-.61
	High	Low	High	-.50
	High	High	Low	-.39
	High	Low	Low	-.29
	Low	High	High	-.29
	Low	Low	High	-.18
	Low	High	Low	-.08
	Low	Low	Low	.03
Secondary	High			-.84
	Low			-.40

The data are from the largest districts, as defined in the text. The impact coefficient indicates the effect on achievement growth, in standard deviations, as contract restrictiveness moves from one standard deviation below its mean to one standard deviation above, given that the stipulated conditions obtain. A “high” value for the relevant condition represents the score of a school at the 75th percentile on that condition, and a “low” score represents the score of a school at the 25th percentile.

References

- Argyris, Laura M., and Daniel I. Rees. 1995. Unionization and School Productivity. *Research in Labor Economics* Vol. 14: 49-68.
- Baumgartner, Frank R., and Beth L. Leech. 1998. Basic Interests. Princeton, N.J.: Princeton University Press.
- Blais, Andre, Donald E. Blake, and Stephane Dion. 1997. Governments, Parties, and Public Sector Employees. Pittsburgh: University of Pittsburgh Press.
- Carpenter, Daniel. 2001. The Forging of Bureaucratic Autonomy. Princeton, NJ: Princeton University Press.
- Ciglar, Alan J., and Burdett A. Loomis. 2006. Interest Group Politics. Washington, D.C.: CQ Press.
- Doucouliafos, Christos, and Patrice Laroche. 2003. What Do Unions Do to Productivity? A Meta-Analysis. Industrial Relations Vol. 42 No.4: 650-691.
- Eberts, Randall W., and Joe A. Stone. 1984. Unions and Public Schools: The Effect of Collective Bargaining on American Education. Lexington, MA: Lexington Books.
-, and Joe A. Stone. 1986. Teachers Unions and the Productivity of Public Schools. *Industrial and Labor Relations Review* Vol. 40: 355-363.
- Epstein, David, and Sharyn O'Halloran. 1999. Delegating Powers. New York: Cambridge University Press.
- Freeman, Richard B. 1986. Unionism Comes to the Public Sector. 24 Journal of Economic Literature: 41-86.
- and Casey Ichniowski, eds., *When Public Sector Workers Unionize*. Chicago: University of Chicago Press.
- Grimes, Paul W., and Charles A. Register. 1990. Teachers Unions and Student Achievement in High School Economics. *Journal of Economic Education* Vol. 21: 297-308.
- Grimshaw, William J. 1979. *Union Rule in the Schools*. Lexington, MA: Lexington Books.
- Hanushek, Eric A. 2002. Publicly Provided Education. In Alan J. Auerbach and Martin Feldstein, eds., Handbook of Public Economics. Amsterdam: North Holland.
-2003. The Failure of Input-Based Schooling. Economic Journal 113: 64-98.
-, John F. Kain, and Steven G. Rivkin. 2004. Why Public Schools Lose Teachers. The Journal of Human Resources 39 no.2 (Spring): 326-354.

Education Working Paper Archive

Hirsch, Barry T. 2004. What Do Unions Do for Economic Performance? *Journal of Labor Research* 25 No. 3: 415-456.

Hoxby, Caroline M. 1996. How Do Teachers Unions Affect Education Production? *Quarterly Journal of Economics* 111: 671-718.

Huber, John D., and Charles R. Shipan. 2002. *Deliberate Discretion? The Institutional Foundations of Bureaucratic Autonomy*. New York: Cambridge University Press.

Johnson, Ronald J., and Gary D. Libecap. 1994. *The Federal Civil Service System and the Problem of Bureaucracy*. Chicago: University of Chicago Press.

Kane, Thomas J., Jonah Rockoff, and Douglas Staiger. 2006. What Does Certification Tell Us about Teacher Effectiveness: Evidence from New York City. NBER Working Paper No. 12155.

Kurth, Michael M. 1987. Teachers Unions and Excellence in Education: An Analysis of the Decline in SAT Scores. *Journal of Labor Research* Vol. 8: 351-367.

Levin, Jessica, Jennifer Mulhern, and Joan Schunck. 2005. "Unintended Consequences: The Case for Reforming the Staffing Rules in Urban Teachers Union Contracts." Report by the New Teacher Project, on the web at www.tntp.org (date of access: December 17, 2005).

Lewin, David, Peter Feuille, Thomas A. Kochan, and John Thomas Delaney, eds.. 1988. *Public Sector Labor Relations*. Lexington, MA: Lexington Books.

Lewis, David. 2003. *Presidents and the Politics of Agency Design*. Stanford, CA: Stanford University Press.

Lieberman, Myron. 1997. *The Teacher Unions*. New York: Free Press.

Loveless, Tom. 2000. *Conflicting Missions: Teachers Unions and Educational Reform*. Washington, D.C.: Brookings Institution.

McCubbins, Mathew D., Roger G. Noll, and Barry R. Weingast. 1987. Administrative Procedures as Instruments of Political Control. *Journal of Law, Economics, and Organization* 3: 243-277.

McDonnell, Lorraine M., and Anthony Pascal. 1979. *Organized Teachers in American Schools*. Santa Monica, CA: Rand Corp.

Milkman, Martin I. Teachers Unions, Productivity, and Minority Student Achievement." *Journal of Labor Research* Vol. 18: 137-150.

Mishel, Lawrence, and Richard Rothstein. 2002. *The Class Size Debate*. Washington, D.C.: Economic Policy Institute.

Moe, Terry M. 2001. Teachers Unions and the Public Schools. In Terry M. Moe, ed., *A Primer on America's Schools*. Stanford, CA: Hoover Press.

..... 2003. Teacher Unions and School Board Elections. In William G. Howell, ed., Beseiged: School Boards and the Future of Education Politics. Washington, D.C.: Brookings Institution.

..... 2006a. Political Control and the Power of the Agent. Journal of Law, Economics, and Organization 22 no. 1 (Spring).

..... 2006b. Union Power and the Education of Children. In Andrew J. Rotherham and Jane Hannaway, eds., Collective Bargaining in Education. Cambridge, MA: Harvard Education Press.

National Commission on Excellence in Education. 1983. A Nation at Risk. Washington, D.C.: Government Printing Office.

Nelson, F. Howard, and Michael Rosen. 1996. Are Teachers unions Hurting American Education? A State-by-State Analysis of the Impact of Collective Bargaining among Teachers on Student Performance. Milwaukee: Institute for Wisconsin's Future.

Peltzman, Sam. 1993. The Political Economy of the Decline of American Public Education. Journal of Law and Economics Vol. 36: 331-370.

Peterson, Paul E. 2003. Our Schools and Our Future. Stanford, CA: Hoover Press.

....., and Martin West. 2003. No Child Left Behind? The Politics and Practice of School Accountability. Washington, D.C.: Brookings Institution.

Rivkin, Steven G., Eric A. Hanushek, and John F. Kain. 2001. Teachers, Schools, and Academic Achievement. Working Paper No. 6691. National Bureau of Economic Research (revised).

Rourke, Francis E. 1984. Bureaucracy, Politics, and Policy. NY: Harper Collins.

Sanders, William, and J. Rivers. 1996. Research Project Report: Cumulative and Residual Effects of Teachers on Future Student Academic Achievement. Knoxville: University of Tennessee, Value Added Research and Assessment Center.

Steelman, Lala Carr, Brian Powell, and Robert M. Carini. 2000. Do Teacher Unions Hinder Educational Performance? Lessons Learned from State SAT and ACT Scores. Harvard Educational Review Vol. 70 (Winter): 437-465.

Troy, Leo. 1994. The New Unionism in the New Society. Fairfax, VA: George Mason University Press.

U.S. Census Bureau, Statistical Abstract of the United States: 2007 (Washington, D.C.: U.S. Census Bureau).